

(12) **United States Patent**
Narayanan et al.

(10) **Patent No.:** US 9,171,041 B1
(45) **Date of Patent:** Oct. 27, 2015

(54) **RLE-AWARE OPTIMIZATION OF SQL QUERIES**

(75) Inventors: **Sivaramakrishnan Narayanan**, Sunnyvale, CA (US); **Florian Michael Waas**, San Francisco, CA (US)

(73) Assignee: **Pivotal Software, Inc.**, Palo Alto, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 31 days.

(21) Appl. No.: **13/248,536**

(22) Filed: **Sep. 29, 2011**

(51) **Int. Cl.**
G06F 17/30 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 17/30463** (2013.01); **G06F 17/30442** (2013.01)

(58) **Field of Classification Search**
CPC G06F 17/30463; G06F 17/30433; G06F 17/30474; G06F 17/30442; G06F 17/30448; G06F 17/30454; G06F 17/3046
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,760,724 B1 * 7/2004 Chakrabarti et al. 1/1
7,167,853 B2 * 1/2007 Zaharioudakis et al. 1/1
7,773,815 B2 * 8/2010 Shen et al. 382/232
8,126,870 B2 * 2/2012 Chowdhuri et al. 707/713

8,171,031 B2 * 5/2012 Tankovich et al. 707/741
8,356,027 B2 * 1/2013 Jaecksch et al. 707/713
2004/0034616 A1 * 2/2004 Witkowski et al. 707/1
2005/0192943 A1 * 9/2005 Siddiqui et al. 707/3
2008/0071818 A1 * 3/2008 Apanowicz et al. 707/101
2010/0088309 A1 * 4/2010 Petculescu et al. 707/714
2011/0131199 A1 * 6/2011 Simon et al. 707/714
2011/0307471 A1 * 12/2011 Sheinin 707/714
2011/0307521 A1 * 12/2011 Slezak et al. 707/800
2012/0054225 A1 * 3/2012 Marwah et al. 707/769
2012/0054236 A1 * 3/2012 Branscome et al. 707/770
2012/0173517 A1 * 7/2012 Lang et al. 707/722

* cited by examiner


Primary Examiner — Leslie Wong

(74) Attorney, Agent, or Firm — Fish & Richardson P.C.

(57) **ABSTRACT**

SQL queries are optimized to operate directly on compressed data (and obtain the correct result) rather than requiring that the data be first decompressed prior to processing a query. Certain characteristic pattern trees are mapped against a logical input query plan that includes certain logical operators such as a DECOMPRESS that precedes a JOIN or a GROUPBY in association with a COUNT to identify instances in the plan that match a characteristic pattern. Upon locating a match, the input query plan is transformed into a logically equivalent plan that operates correctly on compressed data, by analyzing the interplay of the semantics of logical query operations with the compressed data and substituting less costly structures and operations. DECOMPRESS operations are moved to operate subsequent to a JOIN or eliminated altogether, and COUNT operations are replaced by a different operation, such as SUM, that is logically equivalent for compressed data.

16 Claims, 9 Drawing Sheets






Figure 1
(Prior Art)




Figure 2




Figure 3




Figure 4




Figure 5B




Figure 5A




Figure 6




Figure 7




Figure 8




Figure 9A

Figure 9B

1

RLE-AWARE OPTIMIZATION OF SQL QUERIES

BACKGROUND

This invention relates generally to database query optimization, and more generally to optimization of queries for databases having compressed data.

In order to conserve resources and processing time, many databases are compressed at the storage level using, for example, run-length encoding (RLE) compression or other compression techniques. Compression conserves storage space and reduces the number of read requests. However, queries conventionally decompress compressed stored data and operate on uncompressed data. Thus, in order to execute a query on compressed databases, the data must be first decompressed and the query executed multiple times on decompressed data that may be the same. The multiple intermediate results must then be aggregated to obtain an answer. This is inefficient, and results in substantial and costly processing and long overall response times.

It is desirable to provide systems and methods that address these and other known problems of executing queries on compressed data by minimizing computation costs and reducing query response time, and it is to these ends that the invention is directed.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates a conventional shared nothing architecture for a distributed database of the type in which the invention may be employed;

FIG. 2 illustrates a master node of the shared nothing database of FIG. 1 which may incorporate an embodiment of the invention;

FIG. 3 is an example of a logical query plan for a database;

FIG. 4 illustrates the logical query plan of FIG. 3 for use on compressed relations;

FIGS. 5A and 5B illustrate, respectively, an example of a query plan pattern and a transformation of the query plan into a logically equivalent query plan;

FIG. 6 illustrates a query plan pattern that decompresses compressed data and includes a JOIN operator which may be transformed and optimized in accordance with the invention;

FIG. 7 illustrates a transformation and optimization of the query plan of FIG. 6 in accordance with an embodiment of the invention;

FIG. 8 illustrates the transformation of FIG. 7 applied to the query plan of FIG. 4 in accordance with another embodiment of the invention; and

FIGS. 9A and 9B illustrate, respectively, another query plan pattern for use with compressed data and a corresponding transformation of the plan in accordance with an embodiment of the invention, where the transformed plan uses a different query operator to enable operation directly on compressed relations.

DESCRIPTION OF PREFERRED EMBODIMENTS

The invention is particularly well adapted for use with distributed database systems which compresses data at the storage level using RLE-based compression and structured query language (SQL) queries, and will be described in that environment. It will become apparent, however, that this is illustrative of only one utility of the invention and that the

2

invention may be employed with other databases, with other types of compression, and with other query languages.

FIG. 1 illustrates a shared-nothing network architecture of a logical database 100 of the type with which the invention may be employed. The network may include a master node 102 which connects to and manages a plurality of shared-nothing nodes 104-A through 104-N. Each node may comprise a plurality of database segments (database instances) including one or more primary databases and one or more mirror databases. Data may be stored in the segments in compressed form using run length encoding (RLE) compression, for example, and accessed using queries such as SQL queries. Clients 106 may interact with the database through the master node 102.

FIG. 2 illustrates a master node 202 configured to implement operations in accordance with the invention. The master node may comprise a host computer system 210 (which may be a multi-processor system comprising a plurality of CPUs) connected to input/output (I/O) devices 212 by a bus 214. The I/O devices may be standard computer system input and output devices. A network interface circuit 216 may also be connected to bus 214 to allow the master node to operate in a networked environment. Clients 106 may communicate with the distributed database through the master node (see FIG. 1) using either the I/O devices or the network interface circuit. The master node may also have a memory 220 connected to the bus that embodies executable instructions to control the operation of the host computer system of the node and to perform processes in accordance with the invention. Included in memory 220 may be a main memory 222 and a query optimizer 224 comprising instructions that control the CPU to perform query optimization processes in accordance with the invention.

As will be described, the invention affords optimization of queries on databases employing RLE-based compression at the storage level. It provides a framework to exploit RLE compression during query optimization to optimize the queries to minimize computation costs and reduce overall response time. Optimization processes in accordance with embodiments of the invention identify logical query plans, or portions thereof, that bind to certain predetermined patterns and include certain predetermined query operators and/or aggregation operations involving RLE-compressed relations to which logical transformations may be applied that rearrange the query plan and/or use different query operators to produce an optimized logically equivalent plan. The logical query plans are then transformed into the optimized logically equivalent plans. The transformed plans enable optimized queries to operate directly on compressed data and produce correct results without first having to decompress the data. This minimizes computation costs and reduces overall response time by avoiding the necessity of performing multiple query computations on uncompressed data that is the same.

A logical query plan comprises a tree of query language operators that correspond to relational operations such as GET, COMPRESS, DECOMPRESS, JOIN, GROUPBY, and others. The output of each operator is a relation, and each operator produces a set of output columns. For instance, a GET operator corresponds to reading a relation from a storage device and presenting the data in tabular form. This operator has no children. The JOIN operator has two children which correspond to its inputs. It possesses a qualification expression that corresponds to the join condition between two relations. A GROUPBY operator has one child as an input and has a set of grouping columns that should be logically grouped to compute the output relation. The output of the

GROUPBY operator may be associated with aggregate functions such as SUM, MIN and COUNT. A COMPRESS operator has one child, and transforms an uncompressed relation to a compressed one. A DECOMPRESS operator performs the reverse function by transforming a compressed relation into an uncompressed one. For example, a stream of numbers $\langle 5, 5, 6, 6, 6, 6, 1, 1 \rangle$ may be RLE compressed to the form $\langle (5, 2), (6, 4), (1, 2) \rangle$ where second number represents frequency of occurrence of the first element (number). The COMPRESS operation goes from a stream of objects to RLE-compressed form, whereas the DECOMPRESS operation goes in the opposite direction.

FIG. 3 is an example of a logical query plan for a query of the form: SELECT R.r FROM R INNER JOIN T ON (R.r=T.t). R is a relation, e.g., a table, with one column "R.r", where "r" is a column value. Similarly, T is a relation with a column "T.t", where "t" is a column value. The GET (R) operator 302 and the GET (T) operator 304 read values (R.r) 306 and (T.t) 308, respectively. The join operator 310 has a condition "R.r=T.t" 312, that determines whether the values "r" and "t" are equal, and provides the answer to a client 320.

FIG. 4 illustrates a query plan for the same query applied to a compressed database. The RLE-compressed form of relation R may be designated R_c and has two columns ($R_c.r$, $R_c.f$), where the second column corresponds to the frequency of occurrence "f" of a particular "r" value. Similarly, the RLE-compressed form of relation T may be designated T_c with columns ($T_c.t$, $T_c.f$) where the second column indicates the frequency of occurrence of the value "t". In FIG. 4, the operator GET(R_c) 402 obtains the values ($R_c.r$, $R_c.f$) 404 and a DECOMPRESS operator 406 decompress the values to obtain the values "r" 408. Similarly, the operator GET(T_c) 410 obtains the values ($T_c.t$, $T_c.f$) 412, and a DECOMPRESS operator 414 decompress the values to obtain the values "t" 416. A JOIN operator 418 having the condition "R.r=T.t" 420 determines whether the value of "r" and the value of "t" provided by operators 406 and 414, respectively, are equal, and provides the answer to a client 430.

The logical query plan of FIG. 4 compares each decompressed value "r" of R with a given decompressed value "t" of T in the JOIN operation 418, and then repeats this process for the next decompressed value "t". Where the value "r" has a frequency of occurrence of " f_r " and the value "t" has a frequency of occurrence of " f_t ", the query is run f_r times using the same value "r" for each of the f_r occurrences of the same value "r". Thus, the number of times that the query must be run on the same data is f_r multiplied by f_t . If, for example, $f_r=4$ and $f_t=3$, the query must be executed on the same data values $4*3=12$ times, which is costly, time-consuming and inefficient. As will be described, the invention identifies certain patterns in query plans that can be transformed to an optimized plan so that it operates directly on compressed data and provides the correct answer, thus avoiding multiple repeats of a query on data that is the same.

A transformation takes an input query plan and produces a logically equivalent query plan. Every transformation has a precondition that must be satisfied for the transformation to be applicable. The precondition is typically expressed as a pattern tree. If there's a binding (matching) of the input plan to the pattern tree, then the transformation is applicable. FIGS. 5A and 5B illustrate, respectively, a pattern tree and a simple transformation for JOIN commutativity. The transformation shown in FIG. 5B of an input query plan binding to the pattern tree of FIG. 5A merely exchanges the left and right children 502 and 504 of the JOIN operator 506 as shown at 508 and 510 for the JOIN operator 512. As shown in the

figures, both the input plan and the transformation produce the same result 514, which is a requirement for an acceptable transformation.

The invention affords patterns and transformations that enable input query plans, or portions thereof, that bind to certain patterns and which operate on uncompressed data to be transformed so that their transformations operate directly on compressed data to produce the correct answers. While pattern binding and transformations are generally known, they have not previously been applied to compressed data. Known query optimization techniques that are applicable to queries for uncompressed data are very difficult to apply to optimize queries that can be used for compressed data and produce correct answers. The invention identifies those query plans that can be transformed to operate correctly on compressed data by matching the structures and semantics of queries to certain predetermined patterns. The invention then transforms the query plans (and queries) accordingly. Transformations in accordance with the invention take as an input a query plan such as illustrated in FIG. 4 which decompresses relations early in the query so that the query operates on uncompressed data, and transforms the input query plan to one that decompresses higher up (later), or avoids decompression altogether, so that the transformed query plan operates on compressed data. This enables database operations to work on smaller chunks of compressed data, resulting in better performance. This process is illustrated in FIGS. 6 and 7.

FIG. 6 illustrates a pattern for an input plan that is similar to the logical query plan of FIG. 4. As shown, the pattern has DECOMPRESS operators 602 and 604 early in the query that decompress compressed data prior to a JOIN operator 610. This means that the JOIN operates on decompressed data and that the query must be run multiple times on data values that are the same, as previously explained.

FIG. 7 illustrates a transformation of the query plan pattern of FIG. 6 in accordance with an embodiment of the invention, where the DECOMPRESS operator 700 has been moved above the JOIN operator 710 so that DECOMPRESS operates on the results of the JOIN operation. In addition, as indicated at 720, the second output column of the JOIN operator multiplies the frequencies (f_3*f_5) from its left and right children 722, 724, respectively. Notably, the transformed plan of FIG. 7 produces the same answer 730 as the answer 612 produced by the input query plan of FIG. 6, demonstrating that the transformation produces a logically equivalent query that minimizes processing and overall response time to obtain correct answers.

FIG. 8 illustrates another transformation of the logical query plan of FIG. 4 that is substantially similar to the transformation illustrated in FIG. 7, except that the transformation of FIG. 8 eliminates the DECOMPRESS operation altogether from the plan. As shown in FIG. 8, the output of the JOIN operation 810 is supplied directly to the client 812, and the DECOMPRESS operation 820 is moved from the query to the client. This advantageously further reduces the processing that the master node of the database must perform, which improves the database response time to queries. In a large distributed database, moving decompression to a client can result in a substantial improvement in query response time.

FIGS. 9A and 9B illustrate, respectively, another input query plan pattern for use with compressed data, and a corresponding transformation of the input plan in accordance with an embodiment of the invention where the transformed plan uses a different query operation to work directly on compressed relations. The pattern of FIG. 9A applies to an input query on compressed data R_c of the form: SELECT

5

COUNT(*) FROM R_C GROUP BY R_C.r. The pattern has a DECOMPRESS operator 900 that decompresses compressed data 910 before it is operated on by a GROUPBY operator 920 which groups the results of a COUNT operation to provide an answer 940 to the query. The transformation of FIG. 9B changes both the form of the input query as well as its query operators. The query is transformed to: SELECT SUM (R_C.f) FROM R_C GROUP BY R_C.r. As shown, the COUNT operator 930 is replaced with a SUM operator 942. In the transformed plan, the GROUPBY operator 946 produces an answer 948 that is the same as the answer 940 produced by a query operating on decompressed data. This enables eliminating altogether the DECOMPRESS operator 900 from the transformed plan so that it operates on compressed data. This transformation is made possible by the invention because of the discovery that a COUNT operation on uncompressed data in the pattern of FIG. 9A is equivalent to a SUM operation on compressed data in the transformed plan of FIG. 9B.

The foregoing pattern binding and transformation process of the invention may be applied repeatedly to different portions of a more complex query plan to optimize those portions for compressed data and to produce a new overall query plan that is less costly computationally and that has an improved response time.

While the foregoing description has been with reference to particular embodiments of the invention, it will be appreciated by those skilled in the art that modifications to these embodiments may be made without departing from the principles and spirit the invention, the scope of which is defined by the appended claims.

The invention claimed is:

1. A computer-implemented method comprising:
 - receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator followed by another query operator, and (ii) is configured to operate directly on decompressed data;
 - determining that the first pattern of query operators that (i) includes the decompress query operator followed by the other query operator, and (ii) is configured to operate directly on decompressed data, matches a second pattern of one or more query operators that (i) does not include a decompress query operator, and (ii) is configured to operate directly on compressed data;
 - obtaining a transformed query plan by replacing, in the query plan, the first pattern of query operators that (i) includes the decompress operator followed by the other query operator, and (ii) is configured to operate on decompressed data, with the second pattern of one or more query operators that (i) does not include the decompress query operator, and (ii) is configured to operate directly on compressed data; and
 - executing the transformed query plan, including executing the second pattern of one or more query operators that (i) does not include the decompress operator, and (ii) is configured to operate directly on compressed data, on compressed data.
2. The method of claim 1, wherein determining that the first pattern of query operators matches the second pattern of one or query operators, comprises determining, based at least on the (i) structure of the first pattern of query operators and (ii) semantics of the first pattern of query operators, that the first query pattern of query operators matches the second pattern of one or more query operators.
3. The method of claim 1, wherein receiving the query plan that includes the first pattern of query operators that (i) includes the decompress query operator followed by another query operator, and (ii) is configured to operate directly on

6

decompressed data, comprises receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator followed by a JOIN operator, and (ii) is configured to operate directly on decompressed data.

4. The method of claim 3, wherein the one or more query operators which form the second pattern include the JOIN operator.

5. The method of claim 1, wherein receiving the query plan that includes the first pattern of query operators that (i) includes the decompress query operator followed by another query operator, and (ii) is configured to operate directly on decompressed data, comprises receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator followed by a GROUPBY operator, and (ii) is configured to operate directly on decompressed data.

6. The method of claim 5, wherein the first pattern of query operators further includes a COUNT operator that is associated with the GROUPBY operator.

7. The method of claim 6, wherein:

replacing, in the query plan, the first pattern of query operators with the second pattern of one or more query operators comprises replacing, in the query plan, the COUNT operator that is associated with the GROUPBY operator with a SUM operator, and

executing the transformed query plan comprises executing the GROUPBY operator on a sum of compressed data.

8. The method of claim 1, wherein obtaining the transformed query plan comprises obtaining the transformed query plan in response to determining that the first pattern of query operators that (i) includes the decompress query operator followed by the other query operator, and (ii) is configured to operate directly on decompressed data, matches the second pattern of one or more query operators that (i) does not include a decompress query operator, and (ii) is configured to operate directly on compressed data.

9. A non-transitory computer-readable medium storing software comprising instructions executable by one or more computers which, upon such execution, cause the one or more computers to perform operations comprising:

receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator followed by another query operator, and (ii) is configured to operate directly on decompressed data;

determining that the first pattern of query operators that (i) includes the decompress query operator followed by the other query operator, and (ii) is configured to operate directly on decompressed data, matches a second pattern of one or more query operators that (i) does not include a decompress query operator, and (ii) is configured to operate directly on compressed data;

obtaining a transformed query plan by replacing, in the query plan, the first pattern of query operators that (i) includes the decompress operator followed by the other query operator, and (ii) is configured to operate on decompressed data, with the second pattern of one or more query operators that (i) does not include the decompress query operator, and (ii) is configured to operate directly on compressed data; and

executing the transformed query plan, including executing the second pattern of one or more query operators that (i) does not include the decompress operator, and (ii) is configured to operate directly on compressed data, on compressed data.

10. The computer-readable medium of claim 9, wherein determining that the first pattern of query operators matches the second pattern of one or query operators, comprises deter-

mining, based at least on the (i) structure of the first pattern of query operators and (ii) semantics of the first pattern of query operators, that the first query pattern of query operators matches the second pattern of one or more operators.

11. The computer-readable medium of claim 9, wherein receiving the query plan that includes the first pattern of query operators that (i) includes the decompress query operator followed by another query operator, and (ii) is configured to operate directly on decompressed data, comprises receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator followed by a JOIN operator, and (ii) is configured to operate directly on decompressed data.

12. The computer-readable medium of claim 11, wherein the one or more query operators which form the second pattern include the JOIN operator.

13. The computer-readable medium of claim 9, wherein: receiving the query plan that includes the first pattern of query operators that (i) includes the decompress query operator followed by another query operator, and (ii) is configured to operate directly on decompressed data, comprises receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator followed by a GROUPBY operator, and (ii) is configured to operate directly on decompressed data,

the first pattern of query operators further includes a COUNT operator that is associated with the GROUPBY operator,

replacing, in the query plan, the first pattern of query operators with the second pattern of one or more query operators comprises replacing, in the query plan, the COUNT operator that is associated with the GROUPBY operator with a SUM operator, and

executing the transformed query plan comprises executing the GROUPBY operator on a sum of compressed data.

14. The method of claim 1, wherein receiving the query plan that includes the first pattern of query operators that (i) includes the decompress query operator followed by another query operator, and (ii) is configured to operate directly on decompressed data, comprises receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator that is configured to operate on RLE-compressed relations followed by another query operator, and (ii) is configured to operate directly on decompressed data.

15. The computer-readable medium of claim 9, wherein receiving the query plan that includes the first pattern of query operators that (i) includes the decompress query operator followed by another query operator, and (ii) is configured to operate directly on decompressed data, comprises receiving a query plan that includes a first pattern of query operators that (i) includes a decompress query operator that is configured to operate on RLE-compressed relations followed by another query operator, and (ii) is configured to operate directly on decompressed data.

16. The computer-readable medium of claim 9, wherein obtaining the transformed query plan comprises obtaining the transformed query plan in response to determining that the first pattern of query operators that (i) includes the decompress query operator followed by the other query operator, and (ii) is configured to operate directly on decompressed data, matches the second pattern of one or more query operators that (i) does not include a decompress query operator, and (ii) is configured to operate directly on compressed data.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 9,171,041 B1
APPLICATION NO. : 13/248536
DATED : October 27, 2015
INVENTOR(S) : Sivaramakrishnan Narayanan and Florian Michael Waas

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:


In the Claims

Column 5, Lines 58-59, in Claim 2, delete “one or query” and insert -- one or more query --, therefor.

Column 6, Line 67, in Claim 10, delete “one or query” and insert -- one or more query --, therefor.

Column 7, Line 4, in Claim 10, delete “more operators.” and insert -- more query operators. --, therefor.

Signed and Sealed this
Tenth Day of January, 2017

A handwritten signature in black ink, reading "Michelle K. Lee". The signature is fluid and cursive, with the first letters of each name being capitalized and prominent.

Michelle K. Lee
Director of the United States Patent and Trademark Office